

THAT WHICH IS CLAIMED IS:

1. A method of distributing Transmission Control Protocol (TCP) connections to a specific data processing system in a cluster of data processing systems, comprising:

establishing a TCP connection between a client and a first data processing system in the cluster of data processing systems;

obtaining information from the client over the TCP connection through a plurality of request/response communications with the client over the TCP connection;

evaluating the information obtained over the TCP connection to select a target data processing system in the cluster of data processing systems for the TCP connection; and

transferring the TCP connection from the first data processing system to the selected target data processing system so that the transfer of the TCP connection is transparent to the client.

2. A method according to Claim 1, wherein transferring the TCP connection from the first data processing system to the selected target data processing system comprises:

providing connection state information associated with the connection to the selected target data processing system; and

routing subsequent communications associated with the TCP connection to the selected target data processing system to transfer the TCP connection to the selected target data processing system.

3. A method according to Claim 2, wherein evaluating the information obtained over the TCP connection comprises:

providing information received over the TCP connection to an application executing on the first data processing system;

wherein the application executing on the first data processing carries out the following:

evaluating the provided information to select a target instance of an application executing on a target data processing system; and

wherein transferring the TCP connection from the first data processing system to the selected target data processing system further comprises transferring the TCP connection to the selected target instance of the application executing on the target data processing system.

4. The method of Claim 3, further comprising:

notifying the selected target instance of the application executing on the target data processing system of a request to transfer the TCP connection to the selected target instance of the application;

receiving a confirmation of acceptance of the transfer of the TCP connection by the selected target instance of the application; and

wherein transferring the TCP connection to the selected target instance of the application executing on the target data processing system comprises transferring the TCP connection to the selected target instance of the application executing on the target data processing system if the confirmation of acceptance indicates that the selected target instance of the application accepts the transfer of the TCP connection.

5. The method of Claim 3, wherein the target instance of the application executing on the target data processing system comprises a web server.

5 6. The method of Claim 3, wherein the application executing on the first data processing system further carries out the following:

determining application state information based on the provided information and the selected target instance of the application; and

10 wherein transferring the TCP connection from the first data processing system to the selected target data processing system further comprises providing the application state information to the selected target instance of the application executing on the target data processing system.

15 7. The method of Claim 6, wherein the selected target instance of the application executing on the target data processing system carries out the following:

receiving the application state information; and
establishing a state of the target instance of the application based on the received application state information such that the transfer of the TCP connection to the target instance of the application is transparent to the client.

20 8. The method of Claim 6, wherein the application executing on the first data processing system and the target instance of the application executing on the target data processing system are instances of the same application.

9. The method of Claim 6, wherein the application executing on the first data processing system comprises a routing application.

5 10. The method of Claim 1, further comprising transferring the TCP connection from the selected data processing system to a second selected data processing system.

10 11. The method of Claim 1, wherein obtaining information from the client includes peeking at information provided over the TCP connection, and wherein transferring the TCP connection further comprises providing the obtained information to the selected target data processing system.

15 12. A method of transferring a Transmission Control Protocol (TCP) connection to a specific data processing system in a cluster of data processing systems, the cluster of data processing systems having an associated dynamically routable virtual Internet Protocol address (DVIPA), the method comprising:

20 establishing a connection utilizing the DVIPA between a client and a routing application utilizing a routing communication protocol stack at a first data processing system in the cluster of data processing systems;

25 wherein the routing application carries out the following:

30 obtaining information from the client over the connection to the routing application;

selecting a target application for transfer of the connection based on the obtained information; and

35 notifying the routing communication protocol stack of the selected target application;

wherein the routing communication protocol stack carries out the following:

5 sending a connection transfer message to a target communication protocol stack associated with the selected target application, the connection transfer message containing connection state information associated with the connection to the routing application; and

routing subsequent communications over the connection to the target communication protocol stack;

10 wherein the target communication protocol stack carries out the following:

notifying the target application of the transfer of the connection to the target application; and

15 setting a state of a connection to the target application to the state specified by the connection state information associated with the connection to provide a transferred connection to the target application; and

20 wherein the target application carries out communicating with the client utilizing the transferred connection.

13. The method of Claim 12, wherein the routing application further carries out providing application state information to the routing communication protocol stack, the application state information specifying a state of the selected target application based on the information from the client;

25 wherein the routing communication protocol stack further carries out providing the application state information to the target communication protocol stack;

30 wherein the target communication protocol stack further carries out providing the application state information to the target application; and

wherein the target application further carries out resuming communications with the client from the application state specified by the provided application state information utilizing the transferred connection.

14. The method of Claim 12, wherein the target application further carries out sending a response message to the routing application, the response message indicating whether the target application accepts the transfer of the connection.

15. The method of Claim 14, wherein the routing application further carries out closing a socket associated with the connection if the response message from the target application indicates that the target application accepts the transfer of the connection.

16. The method of Claim 14, wherein the routing application further carries out:

selecting a second target application if the response message does not accept the transfer of the connection; and

notifying the routing communication protocol stack of the selection of the second target application so as to initiate transfer of the connection to the second selected target application.

17. The method of Claim 14, wherein the routing application further carries out the step of sending an error message to the client over the connection if the response message indicates that the transfer of the connection is not accepted.

18. The method of Claim 12, wherein the routing application further carries out opening a control socket

to the routing communication protocol stack so as to allow bi-directional communication between the routing communication protocol stack and the routing application.

5 19. The method of Claim 18, wherein the target application further carries out opening a control socket to the target communication protocol stack to allow bi-directional communication between the target application and the target communication protocol stack.

10 20. The method of Claim 18, wherein the routing application is identified to the routing communication protocol stack as a routing application when the routing application opens the control socket.

15 21. The method of Claim 20, wherein the routing communication protocol stack further carries out providing the routing application with an identification of potential target applications listening to the DVIPA by utilizing the control socket.

20 22. The method of Claim 21, wherein the routing communication protocol stack further carries out updating the routing application with identifications of potential target applications listening to the DVIPA by utilizing the control socket.

25 23. The method of Claim 21, wherein the cluster of data processing system includes a plurality of communication protocol stacks and a corresponding plurality of associated applications listening to the DVIPA and wherein the plurality of communication protocol stacks carry out notifying the routing communication protocol stack if a corresponding application is
30 listening to the DVIPA.
35

24. The method of Claim 19, wherein opening a control socket comprises opening a User Datagram Protocol (UDP) socket.

5 25. The method of Claim 12, wherein the cluster of data processing systems comprises a SYSPLEX cluster.

10 26. The method of Claim 11, wherein the routing communication protocol stack further carries out updating a connection routing table associated with the routing communication protocol stack to route communications to the transferred connection to the target communication protocol stack.

15 27. The method of Claim 12, wherein obtaining information from the client comprises peeking information obtained over the connection and wherein the connection transfer message further includes the obtained information.

20 28. A system for distributing Transmission Control Protocol (TCP) connections to a specific data processing system in a cluster of data processing systems, comprising:

25 means for establishing a TCP connection between a client and a first data processing system in the cluster of data processing systems;

30 means for obtaining information through a plurality of request/response communications with the client over the TCP connection;

means for evaluating the information obtained over the TCP connection to select a target data processing system in the cluster of data processing systems for the TCP connection; and

means for transferring the TCP connection from the first data processing system to the selected target data processing system so that the transfer of the TCP connection is transparent to the client.

5
29. A system for transferring a Transmission Control Protocol (TCP) connection to a specific data processing system in a cluster of data processing systems, the cluster of data processing systems having an associated dynamically routable virtual Internet Protocol address (DVIPA), comprising:

10 means for establishing a connection utilizing the DVIPA between a client and a routing application utilizing a routing communication protocol stack at a first data processing system in the cluster of data processing systems;

15 wherein the routing application comprises:

means for obtaining information from the client over the connection to the routing application;

20 means for selecting a target application for transfer of the connection based on the obtained information; and

means for notifying the routing communication protocol stack of the selected target application;

25 wherein the routing communication protocol stack comprises:

30 means for sending a connection transfer message to a target communication protocol stack associated with the selected target application, the connection transfer message containing connection state information associated with the connection to the routing application; and

means for routing subsequent communications over the connection to the target communication protocol stack;

wherein the target communication protocol stack comprises:

means for notifying the target application of the transfer of the connection to the target application; and

means for setting a state of a connection to the target application to the state specified by the connection state information associated with the connection to provide a transferred connection to the target application; and

wherein the target application comprises means for communicating with the client utilizing the transferred connection.

29. A computer program product for distributing Transmission Control Protocol (TCP) connections to a specific data processing system in a cluster of data processing systems, comprising:

a computer-readable storage medium having computer-readable program code embodied therein, the computer-readable program code comprising:

computer-readable program code establishes a TCP connection between a client and a first data processing system in the cluster of data processing systems;

computer-readable program code which obtains information through a plurality of request/response communications with the client over the TCP connection;

computer-readable program code which evaluates the information obtained over the TCP connection to select a target data processing system in the cluster of data processing systems for the TCP connection; and

computer-readable program code which transfers the TCP connection from the first data processing system to the selected target data processing system so that the transfer of the TCP connection is transparent to the client.

30. A computer program product for transferring a Transmission Control Protocol (TCP) connection to a specific data processing system in a cluster of data processing systems, the cluster of data processing systems having an associated dynamically routable virtual Internet Protocol address (DVIPA), comprising:

a computer-readable storage medium having computer-readable program code embodied therein, the computer-readable program code comprising:

computer-readable program code which establishes a connection utilizing the DVIPA between a client and a routing application utilizing a routing communication protocol stack at a first data processing system in the cluster of data processing systems;

wherein the routing application comprises:

computer-readable program code which obtains information from the client over the connection to the routing application;

computer-readable program code which selects a target application for transfer of the connection based on the obtained information; and

computer-readable program code which notifies the routing communication protocol stack of the selected target application;

wherein the routing communication protocol stack comprises:

computer-readable program code which sends a connection transfer message to a target communication protocol stack associated with the selected target application, the connection transfer message containing connection state information associated with the connection to the routing application; and

computer-readable program code which routes subsequent communications over the connection to the target communication protocol stack;

wherein the target communication protocol stack comprises:

computer-readable program code which notifies the target application of the transfer of the connection to the target application; and

computer-readable program code which sets a state of a connection to the target application to the state specified by the connection state information associated with the connection to provide a transferred connection to the target application; and

wherein the target application comprises computer-readable program code which communicates with the client utilizing the transferred connection.